arXiv:1612.08354v1 [cs.CV] 26 Dec 2016

Image-Text Multi-Modal Representation Learning by Adversarial Backpropagation

Gwangbeen Park BOBMARLEY66@KAIST.AC.KR
Woobin Im IWBN@KAIST.AC.KR
Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea

Abstract

We present novel method for image-text multi-modal representation learning. In our knowledge, this work is the first approach of applying adversarial learning concept to multi-modal learning and not exploiting image-text pair information to learn multi-modal feature. We only use category information in contrast with most previous methods using image-text pair information for multi-modal embedding.

In this paper, we show that multi-modal feature can be achieved without image-text pair information and our method makes more similar distribution with image and text in multi-modal feature space than other methods which use image-text pair information. And we show our multi-modal feature has universal semantic information, even though it was trained for category prediction. Our model is end-to-end backpropagation, intuitive and easily extended to other multi-modal learning work.

1. Introduction

Recently, several deep multi-modal learning tasks have emerged. There are image captioning (Vinyals et al., 2015), text conditioned image generation (Reed et al., 2016), object tagging (Karpathy & Fei-Fei, 2015), text to image search (Wang et al., 2016), and so on. For all these works, how to achieve semantic multi-modal representation is the most crucial part.
Therefore, there were several works for multi-modal representation learning (Srivastava & Salakhutdinov, 2012; Frome et al., 2013; Sohn et al., 2014; Wang et al., 2016). And all of these works require image-text pair information. Their assumption is, image-text pair has similar meaning, so if we can embed image-text pair to similar points of multi-modal space, we can achieve semantic multi-modal representation.
But pair information is not always available in several situations. Image and text data usually not exist in pair and if they are not paired, manually pairing them is an impossible task. But tag or category information can exist separately for image and text. And also, does not require paired state and can be manually labeled separately.
And learning multi-modal representation from image-text pair information can be a narrow approach. Because, their training objective focuses on adhering image and text in same image-text pair and doesn't care about adhering image and text, that are semantically similar, but in different pair. So some image and text can have not similar multi-modal feature even though they are semantically similar. In addtion, resolving every pair relations can be a bottleneck with large training dataset.
To deal with above problems, for multi-modal representation learning, we bring concept from ganin's

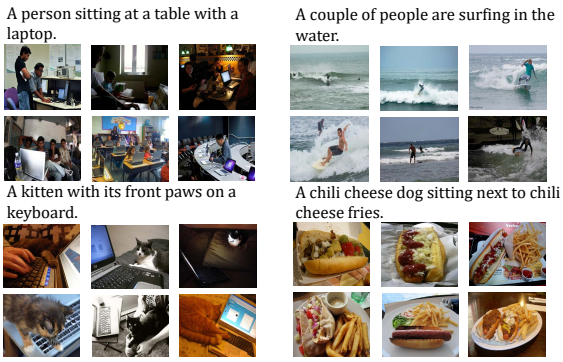
A person sitting at a table with a laptop.
A couple of people are surfing in the water.
A kitten with its front paws on a keyboard.
A chili cheese dog sitting next to chili cheese fries.

Figure 1. Sentence to Image search example from our search system with MS COCO dataset. (upper sentence is query, below images are search result.

work(Ganin & Lempitsky, 2015) which does unsupervised image to image domain adaptation by adversarial backpropagation. They use adversarial learning concept which is inspired by GAN (Generative Adversarial Network)(Goodfellow et al., 2014) to achieve category discriminative and domain invariant feature. We extend this concept to image-text multi-modal representation learning.

We think image and text data are in covariate shift relation. It means, image and text data has same semantic information or labelling function in high level perspective but they have different distribution shape. So we regard, multi-modal representation learning process is adapting image and text distribution to same distribution and retain semantic information at the same time.

In contrast with previous multi-modal representation learning works, we don't exploit image-text pair information and only use category information. Our focus is on achieving category discriminative, domain (image, text) invariant and semantically universal multi-modal representation from image and text.

With above points of view, we did multi-modal embedding with category predictor and domain classifier with gradient reversal layer. We use category predictor for achieving discriminative power of multi-modal feature. And using domain classifier with grdient reversal layer, which makes adversarial relationship with embedding network and domain classifier, for achieving domain (image, text) invariant multi-modal feature. Domain invariant means image and text have same distribution in multi-modal space.

We show that our multi-modal feature distribution is well mixed about domain, which means image and text multi-modal feature's distributions in multi-modal space are similar, and also well distributed by t-SNE(Van Der Maaten, 2014) embedding visualization. And comparison classification performance of multi-modal feature and uni-modal (Image only, Text only) feature shows, there exists small information loss within multi-modal embedding process and still multi-modal feature has category discriminative power even though it is domain invariant feature after multi-modal embedding. And our sentence to image search result (Figure 1) with multi-modal feature shows our multi-modal feature has universal semantic information, which is more than category information. It means, within multi-modal-embedding process, extracted universal information from Word2Vec(Mikolov et al., 2013) and VGG-VeryDeep-16 (Simonyan & Zisserman, 2014) is not removed.

In this paper, we make the following contributions. First, we design novel image-text multi-modal representation learning method which use adversarial learning concept. Second, in our knowledge, this is the first work that doesn't exploit image-text pair information for multi-modal representation learning. Third, we verify image-text multi-modal feature's quality in various perspectives and various methods.

Our approach is much generic as it can be easily used for any different domain (e.g. sound-image, video-text) multi-modal representation learning works with backpropagation only.

2. Related Works

2.1. Multi-Modal Representation

Several works about image-text multi-modal representation learning have been proposed over the recent years. Specific tasks are little bit different for each work, but these works' crucial common part is achieving semantic image-text multi-modal representation from image and text.

Image feature extraction and text feature extraction method are different with each work. But almost they commonly use image-text pair information to learn image-text semantic relation.

Many previous approaches use ranking loss (the training objective is minimizing distance of same image-text pair and maximizing distance of different image-text pair in multi-modal space) for multi-modal embedding. Karpathy's work(Karpathy & Fei-Fei, 2015) use R-CNN(Girshick, 2015) for image feature and BRNN(Schuster & Paliwal, 1997) for text feature and apply ranking loss. And Some approaches ((Frome et al., 2013),(Wang et al., 2016)) use VGG-net for image feature extracting and use neural-language-model for text feature extracting and apply ranking loss or triplet ranking loss.

Some other approaches use deep generative model(Sohn et al., 2014) or DBM (Deep Boltzmann Machine)(Srivastava & Salakhutdinov, 2012) for multi-modal representation learning. In these methods, they intentionally miss one modality feature and generate missed feature from other modality feature to learn relation of different modalities. Therefore they also use image-text pair information and the process is complicate and not intuitive.

2.2. Adversarial Network

Adversarial network concept has started from GAN (Generative Adversarial Network) (Goodfellow et al., 2014). This concept showed great results for several different tasks. For example, DCGAN (Deep Convolutional Generative Adversarial Network) (Radford et al., 2015) drastically improve generated image quality. And text-conditioned DCGAN (Reed et al., 2016)

generate related image from text. Besides image generation, some approach (Ganin & Lempitsky, 2015) apply adversarial learning concept to domain adaptation field with gradient reversal layer. They did domain adaptation from pre-trained image classification network to semantically similar but visually different domain (e.g. edge image, low-resolution image) image target. For this, they set category predictor and domain classifier, which do adversarial learning, so network's feature trained for category discriminative and domain invariant property.

2.3. Domain Adaptation and Covariate Shift

Covariate shift is a primary assumption for domain adaptation field, which assumes that source domain and target domain have same labelling function (same semantic feature or information) but mathematically different distribution form. There was theoretical work about domain adaptation within covariate shift relation source and target domain(Adel & Wong, 2015). And we assume that image and text are also in covariate shift relation. We assume image and text have same semantic information (labeling function) but have different distribution form. So our multi-modal embedding process is adapting those distributions as same and retain semantic information at the same time.

3. Multi-Modal Representation Learning by Adversarial Backpropagation

Now, we detail our proposed novel method and model (Figure 2), multi-modal representation learning by adversarial backpropagation. We use MS COCO dataset (Lin et al., 2014).

3.1. Network Structure

Our network structure (Figure 2) is divided into two parts: feature extraction and multi-modal representation learning. The former part aims at transforming each modality signal into feature. The latter part is devised to embed each feature representation into single (multi-modal) space.

For the representation of visual features, we use VGG16 (Simonyan & Zisserman, 2014) which is pre-trained on ImageNet(Deng et al., 2009). To extract image features, we re-size an image to the size 255×255 and crop 224×224 patches from the four corners and the center. Then the 5 cropped area are flipped to get total 10 patches. We extract fully-connected features (FC7) from each patch, and average them to get a single feature.

To represent sentences, we use Word2Vec (Mikolov et al., 2013), which embeds word into 300-dimensional semantic space. In feature extraction process, words in a sentence are converted into Word2Vec vectors, each of which is a 300-dimensional vector. If a sentence contains N words, we get a feature whose size is $N \times 300$. We add zero padding to the bottom row of the feature to fix its size. Since the maximum length of a sentence in MS COCO dataset (Lin et al., 2014) is 59, we set the feature size to 59×300.

After extracting features from each modality, the multi-modal representation learning process follows.

For an image feature $X \in \mathbb{R}^{4096}$ and a sentence feature $Y \in \mathbb{R}^{59\times300}$, we apply two transformations f^{img} and f^{sent} for images and sentences respectively, to embed two features into a single D-dimensional space. That is, $f^{img}(X) \in \mathbb{R}^D$ and $f^{sent}(Y) \in \mathbb{R}^D$ are satisfied.

For embedding image feature $X \in \mathbb{R}^{4096}$, we use two fully connected layers with ReLU(Nair & Hinton, 2010) activation. Since sentence feature $Y \in \mathbb{R}^{59\times300}$ is 2-dimensional, we apply textCNN (Kim, 2014) to Y, to make it possible for Y to be embedded in the space. At the end of each feature embedding network, we use batch-normalization(Ioffe & Szegedy, 2015) and L2-normalization respectively. And we apply dropout(Srivastava et al., 2014) for all fully-connected layers.

The embedding process is regulated by two components – category predictor and domain classifier with gradient reversal layer, which is a similar concept to that of (Ganin & Lempitsky, 2015). The category predictor regulates the features on the multi-modal space, in such a way that multi-modal features are discriminative enough to be classified into the valid categories. Meanwhile, the domain classifier with gradient reversal layer makes the multi-modal features being invariant to their domain.

3.2. Domain Classifer with Gradient Reversal Layer

We adopt the concept from (Ganin & Lempitsky, 2015).

GRL (Gradient Reversal Layer) is a layer in which backward pass is reversing gradient values. For a layer's input X, the output Y, and the identity matrix I, forward pass is shown on equation 1. In backward pass, for the loss of the network L, the gradient subject to X is shown on equation 2. λ is adaptation factor which is the amount of domain invariance we want to

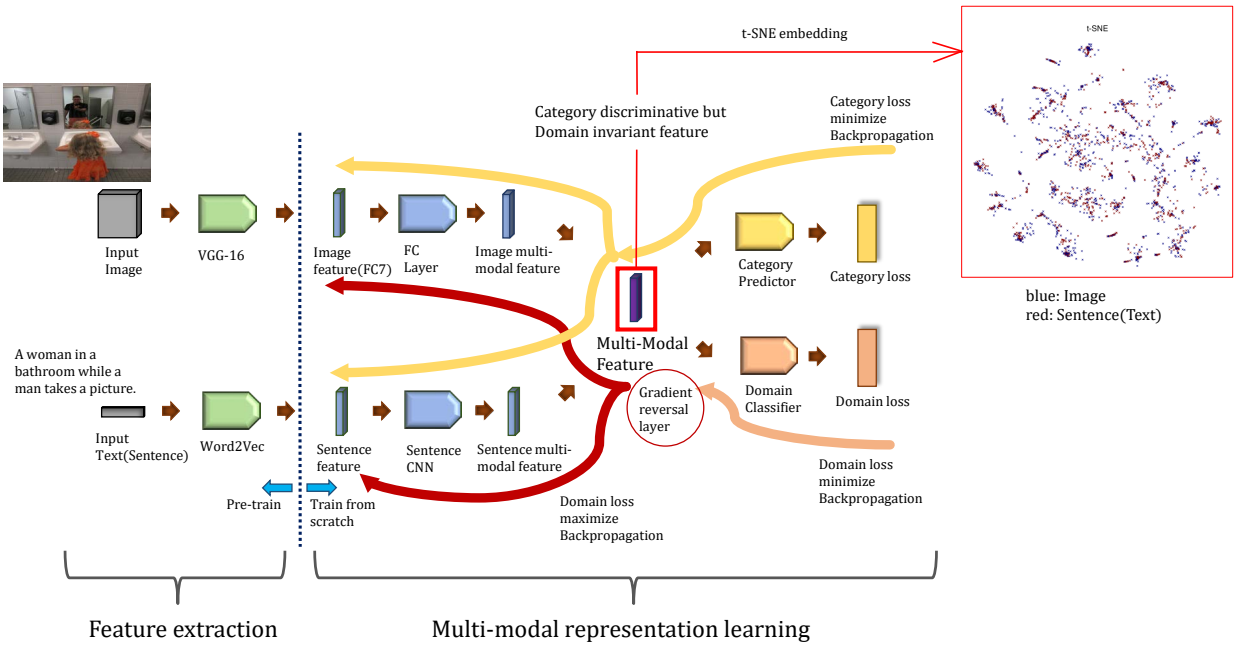

Figure 2. Our adversarial multi-modal embedding network structure. Arrows mean backpropagation of each loss. Feature extraction part is pre-trained by large data set and multi-modal representation learning part is trained with MS COCO dataset's from scratch. Right part of figure is t-SNE embedding of multi-modal feature. (best viewed with color)

achieve at a point of training.

$$Y = IX \tag{1}$$

$$\frac{\partial L}{\partial X} = -\lambda \frac{\partial L}{\partial Y} \tag{2}$$

The domain classifier is a simple neural network that has two fully-connected layers, with the last sigmoid layer that determines the domain of features in the multi-modal embedding space. That is, it is trained in a way that it discriminates the difference between features from two domains.

However, since the GRL reverses the gradient, feature embedding networks are trained to generate features whose domains are difficult to be determined by the domain classifier. This makes adversarial relationship between the embedding network and the domain classifier. Consequently, domain-invariant features can be generated by the multi-modal embedding networks.

3.3. Optimization using Backpropagation

For the calculation of network loss, we sum the losses of two ends – category predictor and domain classifier. We use sigmoid cross entropy loss for the two ends.

For calculating joint gradient of the category predictor and the domain classifier, the two gradients are added, which is shown in the equation below.

$$\frac{\partial L}{\partial Y} = \frac{\partial L_c}{\partial Y} - \lambda \frac{\partial L_d}{\partial Y} \tag{3}$$

where L_c and L_d is the error of category predictor and domain classifier respectively, Y is the output of the last feature embedding layer, and λ is the adaptation factor.

We use Adam optimizer (Kingma & Ba, 2014) for training with relatively small learning rate 10^{-5}. That's because of the empirical difficulty of generating domain-invariant features with regular learning rate.

3.4. Adversarial Learning Procedure

To achieve domain invariant feature, we use domain classifier and gradient reversal layer (Ganin & Lempitsky, 2015). And we should properly schedule λ (adaptation factor) value from 0 to some positive value. Because, at the first stage of training, domain classifier should become smart in advance for adversarial learning process. And with λ value increasing, domain classifying with multi-modal feature become difficult and

domain classifier become smarter to classify it correctly. In our experiment, it turns out that proper λ scheduling is important to achieve domain invariant feature. After exploring many scheduling methods, we find below schedule scheme is optimal, which is exactly the same scheduling as (Ganin & Lempitsky, 2015). In the equation, p is the fraction of current step in max training steps.

$$\lambda = \frac{2}{1 + exp(-10 * p)} - 1 \tag{4}$$

3.5. Normalization

We used batch normalization(Ioffe & Szegedy, 2015) and L2-normalization for normalizing image and text feature distribution just before the multi-modal feature layer. In our experiment, without proper normalization, it seems to be trained well (loss value decreases gently and classification accuracy is fine) but when checking the t-SNE(Van Der Maaten, 2014) embedding and search result, we can recognize that image and text feature distribution is collapsed just for achieving domain invariant feature (collapsed means distance between features going to zero). So proper normalization process is important to achieve domain invariant and also well distributed multi-modal feature.

4. Result

4.1. t-SNE

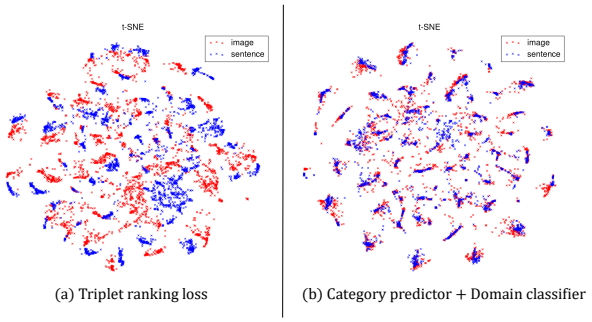

Figure 3. t-SNE embedding with multi-modal feature from image and text. (a) and (b) trained with MS COCO train+validation set . Embedded 5000 data samples come from MS COCO test set. (a): Triplet Ranking Loss (b): Category Predictor + Domain Classifier (our model) (red: image, blue: text) (best viewed with color)

Above Figure 3 is a t-SNE(Van Der Maaten, 2014) embedding result of computed multi-modal features from MS COCO test set's 5000 images and sentences. (a) is result of trained with triplet ranking loss (the training objective is minimizing distance of same image-text pair and maximizing distance of different image-text pair in multi-modal space) which exploits image-text pair relation. For implementing, we consult wang's(Wang et al., 2016) work. We use FV-HGLMM(Klein et al., 2015) for sentence representation, pre-trained VGG16 for image representation and two-branch fully-connected layers for multi-modal embedding, which is same as wang's did. Difference is they use complex data sampling scheme and we use random data sampling at training stage. (b) is result of trained with category predictor and domain classifier which uses our model. In (a), Image and text feature distributions are not well mixed, which means image and text multi-modal feature's distributions in multi-modal space are not similar. Image and text multi-modal features are not overlapped in multi-modal space. It means, semantically similar image and text are not embedded to near points of multi-modal space. But in (b), our result, we can get well mixed with domains image-text multi-modal feature distribution. Image and text are overlapped in multi-modal space and also distributed enough for being discriminated. It means, image and text has similar distribution in multi-modal space. We think difference of (a) and (b) comes from difference of training objective. Because our model(b) trained for hard to classify domain (image, text) of the multi-modal feature and triplet ranking loss(a) is trained for adhering same image-text pair and pushing different image-text pair. And result means triplet ranking loss not adapt image and text to same distribution in multi-modal space. So this result shows our model's training objective is more suitable for learning well mixed with domains and also well distributed multi-modal feature than other methods.

4.2. Classification

Mode	Precision	Recall	F1-score
Image only	0.78	0.81	0.79
Text only	0.78	0.75	0.76
Image+Text(m)	0.72	0.68	0.70

Table 1. category classification result of MS COCO val set with various modes. "Image only" means just use vgg-net and category predictor, "Text only" means just use word2vec, TextCNN and category predictor. So, "Image only" and "Text only" modes don't include domain classifier and gradient reversal layer. "Image+Text(m)" is our multi-modal network model (Figure 2).

Category classification precision/recall result (Table1) shows that after multi-modal embedding, multi-modal feature's (Image+Text(m)) category discriminative power decreases a little, compared to before multi-modal embedding (Image only, Text only), even though it is domain invariant feature. It means our model can adapt image and text feature distribution to multi-modal distribution without large information loss.

4.3. Search

Figure 4. comparison of our sentence to image search result and category based sentence to image search result. We can see our search model not only just care about category information and more universal semantic information even though it just trained for category discriminative multi-modal feature.

We build our sentence to image search system with 40504 number of MS COCO validation set which is never seen at training stage. We train with 82783 images (train set) and test for 40504 images (val set). And simply do k-nearest-neighbor search in multi-modal space with computed multi-modal feature.
Figure 4 shows comparison of our search result and category based search result. You can see the sentence query has more semantic information than its category label. And category based search cannot exploit that semantic information but our search system can exploit that semantic information of sentence query. In figure 4, we can see our search system finds several objects which is not contained in category information but exists in sentence query. In 1st row of figure 4, our search system rightly catches information "woman standing under trees by a field" from sentence even though it was just trained to predict [person, tie] from sentence and image at training time. It means our multi-modal embedding process didn't remove universal information extracted from word2vec and VGG16. And also match image and text semantically relevant feature during multi-modal embedding process.

In 2nd row of figure 4, our search system thinks that most similar image with the query is food image which that category is [spoon, broccoli], which is not overlapped with query's category. But interestingly, in human's semantic perspective, we can recognize they have similar semantic information. ("covered with different vegetables and cheese.")
In Figure 5 (next page), you can see more various search results from our multi-modal search system.

For benchmark of search system, we did recall@K evaluation (table 2, next page) with sentence-to-image and image-to-sentence retrieval. For this, we used karpathy's data split scheme(Karpathy & Fei-Fei, 2015). Compare to state-of-the-art results, our model's performance is relatively low. We think, the major reason is, previous models trained for adhering image-text pair and pushing different image-text pair in multi-modal space and recall@K evaluate query's pair appeared or not in retrieval result. So, even if search result is semantically reasonable (figure 5), if query's pair not appear in retrieval result, recall@K can be low. So we think, this metric is not fully appropriate to assess search quality. But for comparison, we also did recall@K experiment.

5. Conclusion

We have proposed a novel approach for multi-modal representation learning which uses adversarial backpropagation concept. Our method does not require image-text pair information for multi-modal embedding but only uses category label. In contrast, until now almost all other methods exploit image-text pair information to learn semantic relation between image and text feature.
Our work can be easily extended to other multi-modal representation learning (e.g. sound-image, sound-text, video-text). So our method's future work will be extending this method to other multi-modal case.

References

Adel, Tameem and Wong, Alexander. A probabilistic covariate shift assumption for domain adaptation. In *AAAI*, pp. 2476–2482, 2015.

Deng, Jia, Dong, Wei, Socher, Richard, Li, Li-Jia, Li, Kai, and Fei-Fei, Li. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255. IEEE, 2009.

Frome, Andrea, Corrado, Greg S, Shlens, Jon, Bengio, Samy, Dean, Jeff, Mikolov, Tomas, et al. Devise: A

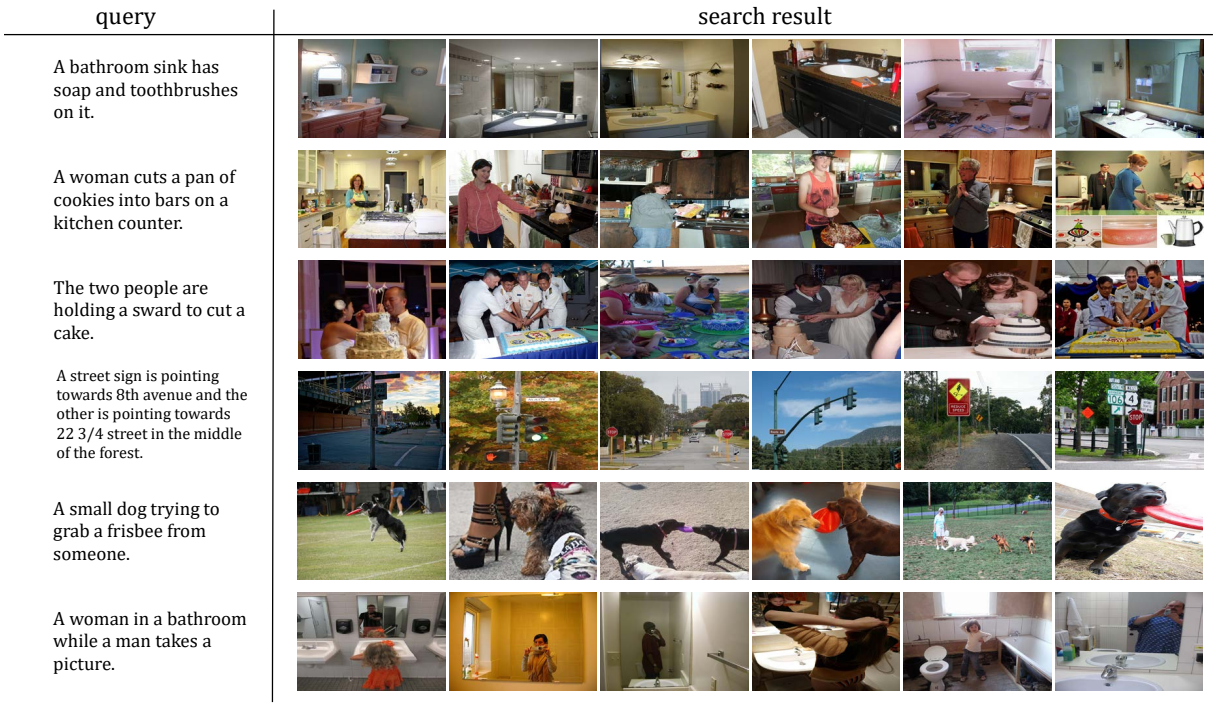
query
search result
A bathroom sink has soap and toothbrushes on it.
A woman cuts a pan of cookies into bars on a kitchen counter.
The two people are holding a sward to cut a cake.
A street sign is pointing towards 8th avenue and the other is pointing towards 22 3/4 street in the middle of the forest.
A small dog trying to grab a frisbee from someone.
A woman in a bathroom while a man takes a picture.

Figure 5. sentence to image search result (top 6) example from 40504 number of MS COCO validation set.

	Model	Image-to-Sentence			Sentence-to-Image		
		R@1	R@5	R@10	R@1	R@5	R@10
Using image-text pair information	DVSA(Karpathy & Fei-Fei, 2015)	38.4	69.9	80.5	27.4	60.2	74.8
	m-RNN-vgg(Mao et al., 2014)	41.0	73.0	83.5	29.0	42.2	77.0
	mCNN (ensemble)(Ma et al., 2015)	42.8	73.1	84.1	32.6	68.6	82.8
	structure preserving(Wang et al., 2016)	50.1	79.7	89.2	39.6	75.2	86.9
Using category information	**our model (TextCNN)**	13.6	39.6	54.6	10.3	35.5	55.5
	our model (FV-HGLMM)	14.3	40.5	55.8	12.7	39.0	57.2

Table 2. Bidirectional retrieval recall@K results on MS COCO 1000-image test set. our model (TextCNN) and our model (FV-HGLMM) use TextCNN and FV-HGLMM for text representation respectively.

deep visual-semantic embedding model. In *Advances in neural information processing systems*, pp. 2121–2129, 2013.

Ganin, Yaroslav and Lempitsky, Victor. Unsupervised domain adaptation by backpropagation. In Blei, David and Bach, Francis (eds.), *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 1180–1189. JMLR Workshop and Conference Proceedings, 2015. URL http://jmlr.org/proceedings/papers/v37/ganin15.pdf.

Girshick, Ross. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1440–1448, 2015.

Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.

Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL http://arxiv.org/abs/1502.03167.

Karpathy, Andrej and Fei-Fei, Li. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3128–3137, 2015.

Kim, Yoon. Convolutional neural networks for sentence classification. In Moschitti, Alessandro, Pang, Bo, and Daelemans, Walter (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pp. 1746–1751. ACL, 2014. ISBN 978-1-937284-96-1. URL http://aclweb.org/anthology/D/D14/D14-1181.pdf.

Kingma, Diederik P. and Ba, Jimmy. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL http://arxiv.org/abs/1412.6980.

Klein, Benjamin, Lev, Guy, Sadeh, Gil, and Wolf, Lior. Associating neural word embeddings with deep image representations using fisher vectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4437–4446, 2015.

Lin, Tsung-Yi, Maire, Michael, Belongie, Serge, Hays, James, Perona, Pietro, Ramanan, Deva, Dollár, Piotr, and Zitnick, C Lawrence. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pp. 740–755. Springer, 2014.

Ma, Lin, Lu, Zhengdong, Shang, Lifeng, and Li, Hang. Multimodal convolutional neural networks for matching image and sentence. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2623–2631, 2015.

Mao, Junhua, Xu, Wei, Yang, Yi, Wang, Jiang, Huang, Zhiheng, and Yuille, Alan. Deep captioning with multimodal recurrent neural networks (m-rnn). *arXiv preprint arXiv:1412.6632*, 2014.

Mikolov, Tomas, Sutskever, Ilya, Chen, Kai, Corrado, Greg S, and Dean, Jeff. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119, 2013.

Nair, Vinod and Hinton, Geoffrey E. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 807–814, 2010.

Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. *CoRR*, abs/1511.06434, 2015. URL http://arxiv.org/abs/1511.06434.

Reed, Scott E., Akata, Zeynep, Yan, Xinchen, Logeswaran, Lajanugen, Schiele, Bernt, and Lee, Honglak. Generative adversarial text to image synthesis. In Balcan, Maria-Florina and Weinberger, Kilian Q. (eds.), *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1060–1069. JMLR.org, 2016. URL http://jmlr.org/proceedings/papers/v48/reed16.html.

Schuster, Mike and Paliwal, Kuldip K. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681, 1997.

Simonyan, Karen and Zisserman, Andrew. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. URL http://arxiv.org/abs/1409.1556.

Sohn, Kihyuk, Shang, Wenling, and Lee, Honglak. Improved multimodal deep learning with variation of information. In *Advances in Neural Information Processing Systems*, pp. 2141–2149, 2014.

Srivastava, Nitish and Salakhutdinov, Ruslan R. Multimodal learning with deep boltzmann machines. In *Advances in neural information processing systems*, pp. 2222–2230, 2012.

Srivastava, Nitish, Hinton, Geoffrey E, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

Van Der Maaten, Laurens. Accelerating t-sne using tree-based algorithms. *Journal of machine learning research*, 15(1):3221–3245, 2014.

Vinyals, Oriol, Toshev, Alexander, Bengio, Samy, and Erhan, Dumitru. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3156–3164, 2015.

Wang, Liwei, Li, Yin, and Lazebnik, Svetlana. Learning deep structure-preserving image-text embeddings. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.